Journal of
Statistics Education

The Not-so-Random Drunkard's Walk

George Ehrhardt
Appalachian State University

Key Words: Data Set and Story**;** Random Walk; Probability; Test of Sample Proportion

Abstract

This dataset contains the results of a quasi-experiment, testing Karl Pearson's "drunkard's walk" analogy for an abstract random walk. Inspired by the alternate hypothesis that drunkards stumble to the side of their dominant hand, it includes data on intoxicated test subjects walking a 10' line. Variables include: the direction to which subjects first stumbled, the side of the line on which they ended up, and a record of subjects' dominant hand. In addition to enhancing a discussion of the random walk itself, instructors may find this dataset useful for teaching tests of sample proportions.

1. Introduction

In 1905, Karl Pearson posed a question in the journal *Nature*, asking for a mathematical solution to predicting how far from his initial location a man would travel if he walks one yard in a straight line, then "turns through any angle whatever," walks another yard, turns, and so on *n* times. After reading the responses his question provoked, Pearson (1905) concluded:

> "In open country, the most probable place to find a drunken man who is at all capable of keeping on his feet is somewhere near his starting point!" (p.342)

Thus was born the "drunkard's walk." Since Pearson wrote those words, the term "drunkard's walk" has sat, perhaps unsteadily, alongside its synonym "random walk" in the study of probability.

Predictably catchy, the term often appears in material for a general audience, from Leonard

Mlodinow's (2008) book *The Drunkard's Walk: How Randomness Rules Our Lives*, a mathematical revision of Sherlock Holmes in Colin Bruce's (2002) *Conned Again, Watson*, to inclusion in Wikipedia (2013).

Mlodinow shows how this process resembling a drunkard stumbling home gives insight into many path-dependent phenomena, like the stock market, music sales, and nuclear reactor safety. Given the connection between drinking and gambling, it comes as no surprise to see it used in calculating the number of times a deck of cards must be shuffled in Kaplan and Kaplan's (2007) *Chances Are* and labeled the "gambler's ruin problem" in Bailey's (1964) *Elements of Stochastic Processes.*

For the college statistics instructor, this term's appeal is a double-edged sword. On the one hand, it illustrates the important concept of a "random walk" in a way that students can intuitively understand. Stumbling home drunk is unfortunately a familiar experience for many of today's college students. Numerous studies have documented the prevalence of excessive drinking on college campuses, along with many consequences. According to studies supported by the National Institute of Alcohol Abuse and Alcoholism (see for example Hingson, Edwards, Heeren, and Rosenbloom (2009), http://www.ncbi.nlm.nih.gov/pmc/articles/PMC2701090/), these include an estimated 1,825 annual deaths of college students from alcohol-related unintentional injuries, 599,000 unintentional injuries, and 97,000 victims of alcohol-related sexual assault. The use of this dataset should not encourage or condone excess drinking. While acknowledging the danger, however, an instructor can use the drunkard's walk in a variety of settings.

One possible use is illustrating how the summation of many random factors gives rise to approximations to the normal distribution. In his book, Bruce (2002) gives an entertaining version of the explanation that involves a drunken sailor stumbling down a pier. He simplifies Pearson's story by assuming that drunken stumbles represent side-steps, sliding laterally while moving generally forward. Sherlock Holmes then explains how the sailor is most likely to end up in the middle, with each possible endpoint having proportionally decreasing probability the farther it is from the center. Assuming equal chances of stumbling left and right with each step, making a histogram of the number of routes to each endpoint produces a discrete distribution that is approximately normal, and students can use this to match the low probability of falling off the pier with the low probability of outcomes in the tails of a normal distribution.

> *Helpful Hint: Because modern statistical software can calculate exact binomial probabilities, the normal approximation is no longer needed for such calculations. However, the convergence due to the Central Limit Theorem is an important concept that is well illustrated by this example.*

While teaching this story one day, a student challenged its factual accuracy. She speculated that drunks tend to stumble in the direction of their dominant side, so a right-handed person would be more likely to stumble right, and a left-handed person stumble left. Others disagreed. To resolve the controversy, our class decided to answer the question empirically, as described in the next section.

> *Helpful Hint: our class discussion used the term "footedness," but we tested the student's claim by asking what hand the subjects wrote with.   This gap between her theory (involving the vague idea of dominant side) and our test (asking a precise question) could spur a discussion of the concept/variable distinction.*

2. Study Design and Data

In an idealized world, testing the student's hypothesis would involve randomly selecting subjects, serving them enough alcohol to get drunk, and measuring their predilection for stumbling. Given the potential consequences, this would be unethical at best, and dangerous at worst. A more appropriate alternative would be to observe subjects who choose to drink without instruction, but this limits the researcher's ability to control the sample.

> *Helpful Hint: This is a common situation in the social sciences. While the following discussion may seem irrelevant to instructors who solely teach statistical analysis, ethical issues are an important issue for students to consider. The American Statistical Association Ethical Guidelines for Statistical Practice, 1999,* http://www.amstat.org/about/ethicalguidelines.cfm *provide a useful introduction.*

Campbell and Stanley (1963) refer to situations like this type of observational study as quasi-experiments, where the researchers have some control over the experimental intervention. In practice, this meant that the design of the study called for student researchers to attend parties on campus, find people who self-identified as drunk, ask them whether they were right-handed or left-handed, how many drinks they'd had, and note their gender (the observational *quasi-*component). When they had recorded that information, they asked the subjects to walk a 10' line taped on the floor (the *experimental* intervention component). As the subjects walked the line, the student observers marked down the first error--was their first stumble to the right or left?

This side-stumbling walk (as modeled in Bruce's story) is slightly different from Pearson's original representation, but it more accurately reflects how drunkards walk. When the subject reached the end, they marked down whether the subject was to the left or right of their original position on the line. To ensure a large enough sample of left-handed subjects, student researchers consciously sought out left-handed drinkers until they had equal numbers of each. To minimize measurement error, student researchers agreed not to drink alcohol until after they had finished their observations.

If Pearson's analogy of a drunkard in a field ending up near where he started is correct, then we should see no pattern in test subjects' walks; there should be no relationship between their dominant hand and the direction in which they stumble first (or last).

On the other hand, if subjects tend to stumble in the direction of their dominant hand, the study would suggest that the drunken walk is not actually random. The Appalachian State University Institutional Review Board approved this research protocol.

3. Description of the Data

The resulting dataset of drunkards' walks, *Drunkwalks.txt*, is available the JSE archives. It contains a total of 151 complete observations: 75 right-handed subjects and 76 left-handed subjects.

After an initial case ID variable, the next two columns contain possible control variables: how many drinks the subjects' recalled consuming prior to testing and their sex. The next column indicates their dominant hand. The subsequent two columns record the outcomes of the subjects' walk: one the direction of their first stumble and the other whether they end up to the left or right of the original line. The data does not record how far off individuals strayed.   The final column is a calculated variable, with a value of 1 if the dominant hand variable and first stumble match (e.g. a right-hander stumbling right) and zero if they do not (e.g. a left-hander stumbling right). More detail on the variables and coding is available in the documentation file *drunkwalk_coding.txt*.   A numerical summary of the data is provided in Table 1.

Table 1: Numerical Summary of the Data

	Average	**Standard Deviation**	**Min**	**Max**
# of Drinks:	7.3	2.7	2	14
	Male	**Female**		
Sex	68	83		
	Right	**Left**		
Handedness	75	76		
	Right	**Left**		
First step off	75	76		
Ending Point	67	84		
	Yes	**No**		
Hand/Stumble match	107	44		

Helpful hint: Only one subject in the study successfully walked the line, even those at the bottom end of the drinks distribution. In describing this dataset to students, the instructor may wish to point out how few drinks it took before students (effectively) failed a field sobriety test. Even "just a few drinks" is enough to land them in jail with a DUI.

4. Testing Claims about the Drunken Walk

This dataset is suitable for many different statistical tests. The 2x2 structure comparing dominant hand's side to first step off, or to ending position, lends itself to chi-squared tests or

measures of association with nominal data. Those analyses would examine hypotheses about whether there is a connection between a subject's handedness and their "direction of stumble."

4.1 Testing the initial claim of non-randomness

The particular claim that prompted this research, however, did not only assert a correlation between handedness and stumbling, but a particular relationship--that drunkards stumble towards the side of their dominant hand. We do this with a test of sample proportions, in which stumbling towards the subject's dominant hand is coded as a success, and stumbling toward the non-dominant hand is coded as a failure. This test, ubiquitous in introductory statistics classes, tests whether the population divides on a dichotomous variable by a hypothesized percent.

> *Helpful Hint: the creation of a new variable to measure whether subjects stumble to the side of their dominant hand (rather than thinking in terms of left->left and right->right in a divided dataset) is not intuitive to many students. This may be an opportunity to discuss the importance of data preparation in facilitating analysis.*

If Pearson is correct, then no matter what a subjects' handedness may be, they have a 50% chance of stumbling to that side, producing the following hypothesis:

$$H_a: \pi_m = 0.5$$

where π_m is the proportion of drunkards who stumble towards their dominant hand's side. The claim that "people tend to stumble towards their dominant hand's side," on the other hand, can be formalized as:

$$H_a: \pi_m > 0.5$$

As stated, this hypothesis suggests a one-sided test, because the claim makes an affirmative assertion about the direction of stumbling; finding that the proportion equals or is less than 0.5 would both lead us to reject the student's hypothesis. In practice, however, it could be performed as a two-sided test ($H_b: \pi_m \neq 0.5$), interpreting the results similarly as long as the observed probability is greater than 0.5 (results for both one- and two-sided tests are included below).

There are two criteria that must be met to use this test: the expected number of successes and failures must be larger than 10, and the data must be taken from a random sample. In this case, the first criteria is met (since $\pi_m = 0.5$ and *n*=151) but the second is not.

> *Alternative Application: given this issue with the test's assumptions, instructors could also use an exact binomial test on the data, or utilize a resampling based technique (*[Chihara and Hesterberg 2012](#)*). The choice should be based on the goals and structure of the course and is beyond the scope of this paper.*

The convenience stratified sample of drunk students, where left-handers were oversampled is not a simple random sample. Nevertheless, use of non-random samples is common in some fields,

with an appropriate disclaimer (Garson 2012). That seems especially true in situations like this one, where we cannot ethically choose a random sample and supply them with alcohol.

> *Helpful Hint: one way to deal with this is as an opportunity to point out to students how frequently true non-random samples are used, and how these may not be generalizable to the population of interest. This also suggests that many of the studies they read may have more uncertainty due to non-sampling errors than the reported significance levels that account for sampling errors.*

Students recorded 107 out of 151 subjects stumbling to their dominant hand's side, a proportion of 0.71. Running the test in STATA produced a z-score of 5.13 and a one-tailed p-value of <0.0001; the 95% confidence interval is 0.64 to 0.78. The two-tailed p-value is also <0.0001 These results suggest (with the appropriate caveat for a non-random sample) that there is a pattern in how drunkards stumble—towards the side of their dominant hand.

4.2 Testing whether the amount a subject drinks affects their direction of stumble

In the course of our discussion, one student wondered whether it matters how many drinks a subject consumed prior to walking the line. This too can be determined from the data at hand with another test for a population proportion, using the **drinks** variable to divide our sample.

For simplicity, we divided our sample into two samples; since the mean value for **drinks** was 7.3, we used 7 as the cut point. This produced two samples, described in Table 2: numerical summary of each sample.

Table 2: Numerical summary of drink variable for each sample

	"Fewer drinks" (≤7) Sample (n=86)	"More Drinks" (>7) Sample (n=65)
Minimum	2	8
Mean	5.3	9.8
Maximum	7	14
Standard deviation	1.45	1.65
proportion stumbling to dominant side	0.67	0.75

Unlike the first case, here we are comparing two samples to test whether their proportion of matches (between stumble side and dominant hand's side) is the same or different, with the same caveats about non-random sampling as discussed above. The competing hypotheses for this two-sample test of proportion are:

$$H_0: \pi_m(fewer\ drinks) = \pi_m(more\ drinks)$$
$$H_a: \pi_m(fewer\ drinks) \neq \pi_m(more\ drinks)$$

Since the student did not hypothesize whether more alcohol would make one more or less likely to stumble towards one's dominant hand's side, we conducted this as a two-tailed test.

> *Helpful Hint: This could spark a discussion about whether a two-sided or one-sided test is appropriate in this setting.*

The sample that had fewer drinks reported 58 of 86 matches, for a proportion of 67%, while the sample with more drinks reported 49 of 65 matches, for a proportion of 75%. The sample with more drinks under their belt did tend to stumble towards the side of their dominant hand. Running this test however, revealed that the difference was not large enough to be confident in inferring differences in the population; the z-score= -1.06 and the two-tailed p-value = 0.29 (the 95% confidence interval for the difference in proportions ranged from -0.22 - 0.06).

5. Conclusion

Karl Pearson's analogy of a drunkard stumbling home to describe the mathematical process of a random walk is an enduring and compelling image in the study of probability. Unfortunately, it is also a regular part of life for many college students. While we must be mindful of the danger that poses, the drunkard's walk analogy offers a teaching tool that raises student interest and presents important concepts in a context the students can easily understand.

Useful though the analogy may be, empirically it leaves something to be desired. In response to student doubt about whether drunks actually stumble randomly, a class tested the proposition among their peers. In a quasi-experiment asking inebriated college students to walk a 10' line, they found that contrary to Pearson's supposition, drunkards tend to stumble towards the side of their dominant hand. The strength of this result does not appear to depend on how much the subject has drunk, beyond the initial threshold of being unable to walk the 10' without stumbling. The results of this quasi-experiment can be found in the JSE archives in the file *drunkardswalks.dat*, along with a documentary file explaining its contents.

References

Bailey, N. (1964), *Elements of Stochastic Processes,* New York: Wiley and Sons.

Bruce, C. (2002), *Conned Again,* Watson. Cambridge, MA: Perseus Books.

Campbell, D. and Stanley, J. (1963), *Experimental and Quasi-Experimental Designs for Research*. Boston: Houghton Mifflin.

Chihara, L. M. and Hesterberg, T. C. (2012), *Mathematical Statistics with Resampling and R*. New York: John Wiley and Sons.

Garson, G. David (2012), *Sampling*. Statistical Associates Publishing.

Hingson, R.W., Edwards, E.M., Heeren, T. and Rosenbloom, D. (2009), Age of drinking onset and injuries, motor vehicle crashes, and physical fights after drinking and when not drinking., *Alcohol Clinical and Experimental Research*, 33(5):783-790.

Kaplan M. and E. Kaplan (2007), *Chances are: Adventures in Probability* New York: Penguin Books.

Mlodinow, L. (2008), *The Drunkard's Walk: How Randomness Rules our* Lives. New York, New York: Pantheon Books.

Pearson, K. (1905), "The Problem of the Random Walk" *Nature* no. 1867 v. 72, pp. 294 and 342. Wikipedia contributors, "Random walk," *Wikipedia, The Free Encyclopedia*, http://en.wikipedia.org/w/index.php?title=Random_walk&oldid=529945744 (accessed January 15, 2013).

George Ehrhardt
Appalachian State University